Model Interpolation with Trans-dimensional Random Field Language Models for Speech Recognition

Bin Wang[1], *Zhijian Ou*[1], *Yong He*[2], *Akinori Kawamura*[2]

[1]Department of Electronic Engineering, Tsinghua University, Beijing 100084, China
[2]Toshiba Corporation
contact email: ozj@tsinghua.edu.cn

arXiv:1603.09170v5 [cs.CL] 20 Aug 2016

Abstract

The dominant language models (LMs) such as n-gram and neural network (NN) models represent sentence probabilities in terms of conditionals. In contrast, a new trans-dimensional random field (TRF) LM has been recently introduced to show superior performances, where the whole sentence is modeled as a random field. In this paper, we examine how the TRF models can be interpolated with the NN models, and obtain 12.1% and 17.9% relative error rate reductions over 6-gram LMs for English and Chinese speech recognition respectively through log-linear combination.

Index Terms: Language modeling, random fields, stochastic approximation

1. Introduction

Language modeling (LM) involves determining the joint probability of words in a sentence. The conditional approach is dominant, representing the joint probability in terms of conditionals. Examples include n-gram LMs [1] and neural network (NN) LMs [2, 3]. Recently a new trans-dimensional random field (TRF) LM [4] is introduced, where the whole sentence is modeled as a random field. As the random field approach avoids local normalization which is required in the conditional approach, it is computationally more efficient in computing sentence probabilities and has the potential advantage of being able to flexibly integrate a richer set of features. An effective training algorithm using joint stochastic approximation (SA) and trans-dimensional mixture sampling is developed in [4]. It is found that the TRF models significantly outperformed the modified Kneser-Ney (KN) smoothed 4-gram LM with 9.1% relative reduction in speech recognition word error rates (WERs) and performed slightly better than the recurrent neural network LMs but with 200x faster speed in re-scoring n-best lists of hypothesized sentences. To our knowledge, this result represents the first strong empirical evidence supporting the power of using the whole-sentence random field approach for LMs [5].

In this paper, we examine how the TRF models can be interpolated with NN models to further improve the performance. Speech recognition experiments are conducted for both English and Chinese. Three kinds of interpolation schemes, namely linear interpolation (word-level or sentence-level) and log-linear interpolation, are evaluated for model combinations between KN n-gram LMs, NN LMs and our TRF LMs. Through log-linearly combining TRF LMs with NN LMs, we obtain 12.1% and 17.9% relative error rate reductions over the KN 6-gram LMs for English and Chinese speech recognition respectively.

This work is supported by NSFC grant 61473168, Tsinghua Initiative 20121088069 and Toshiba Corporation.

2. Background of TRF Model Definition and Training

Throughout, we denote by $x^l = (x_1, \ldots, x_l)$ a sentence (i.e., word sequence) of length l ranging from 1 to m. Each element of x^l corresponds to a single word. D denotes the whole training corpus and D_l denotes the collection of length l in the training corpus. n_l denotes the size of D_l and $n = \sum_{l=1}^{m} n_l$.

As defined in [4], a trans-dimensional random field model represents the joint probability of the pair (l, x^l) as

$$p(l, x^l; \lambda) = \frac{n_l/n}{Z_l(\lambda)} e^{\lambda^T f(x^l)}, \tag{1}$$

where n_l/n is the empirical probability of length l. $f(x^l) = (f_1(x^l), \ldots f_d(x^l))^T$ is the feature vector, which is usually defined to be position-independent and length-independent, e.g. the n-grams. d is the dimension of the feature vector $f(x)$. λ is the corresponding parameter vector of $f(x^l)$. $Z_l(\lambda) = \sum_{x^l} e^{\lambda^T f(x^l)}$ is the normalization constant of length l. By making explicit the role of length in model definition, it is clear that the model in (1) is a mixture of random fields on sentences of different lengths (namely on subspaces of different dimensions), and hence will be called a trans-dimensional random field (TRF).

In the joint SA training algorithm [4], another form of mixture distribution is defined as follows:

$$p(l, x^l; \lambda, \zeta) = \frac{n_l/n}{Z_1(\lambda) e^{\zeta_l}} e^{\lambda^T f(x^l)} \tag{2}$$

where $\zeta = \{\zeta_1, \ldots, \zeta_m\}$ with $\zeta_1 = 0$ and ζ_l is the hypothesized value of the log ratio of $Z_l(\lambda)$ with respect to $Z_1(\lambda)$, namely $\log \frac{Z_l(\lambda)}{Z_1(\lambda)}$. $Z_1(\lambda)$ is chosen as the reference value and can be calculated exactly. An important observation is that if and only if ζ were equal to the true log ratios, then the marginal probability of length l under distribution (2) equals to n_l/n. This property is then used to construct the joint SA algorithm, which jointly estimates the model parameters and normalization constants. More recent development on TRF models and related algorithms, see [6].

3. Experiments

We examine how the TRF models can be interpolated with NN models to further improve the performance. Linear combination can be done at either word-level (W) or sentence-level (S). In contrast, log-linear combination has no such differentiation. The three schemes are detailed in Tab.1. As TRF models only output sentence probabilities, the "W" scheme is not applicable in combining TRF models with other models.

W	$s(x^l) = \prod_{i=1}^{l} [\alpha p_1(x_i\|h_i) + (1-\alpha)p_2(x_i\|h_i)]$
S	$s(x^l) = \alpha p_1(x^l) + (1-\alpha)p_2(x^l)$
Log	$s(x^l) = \exp(\alpha \log p_1(x^l) + (1-\alpha)\log p_2(x^l))$

Table 1: Model interpolation schemes. $s(x^l)$ is the combined score of sentence x^l. $p_1(x_i|h_i)$ and $p_2(x_i|h_i)$ are the conditional probabilities of x_i given history h_i estimated by two LMs; $p_1(x^l)$ and $p_2(x^l)$ are the joint probabilities of the whole sentence x^l estimated by two LMs. $0 < \alpha < 1$ is the interpolation weight which is tuned on the development set.

Model	WER	PPL	#feat (M)
KN5 (cutoff=00000)	8.61	293.94	2.3
KN6 (cutoff=000000)	8.61	293.67	3.0
RNN [3]	7.96	256.15	5.1
TRFs (200 classes)			
w+c+ws+cs+cpw	7.92	∼257	6.4
w+c+ws+cs+wsh+csh	7.94	∼259	5.9
w+c+ws+cs+wsh+csh+tied*	**7.85**	∼261	7.7
The above results are reproduced from [6].			
Model Combination			
RNN + KN5 (W)	7.85		
RNN + KN5 (S)	7.96		
RNN + KN5 (Log)	7.96		
TRF*+ RNN (S)	**7.51**		
TRF*+ RNN (Log)	**7.57**		

Table 2: The WERs and PPLs on the WSJ'92 test data. "#feat" denotes the feature size (million). "(W),(S),(Log)" denote different model combination schemes as defined in Tab.1. "TRF*" denotes the TRF with features "w+c+ws+cs+wsh+csh+tied". "tied" represents the tied long-skip-bigram features introduced in [6], in which the skip-bigrams with skipping distances from 6 to 9 share the same parameter.

3.1. English speech recognition on WSJ0 dataset

In this section, speech recognition and 1000-best list rescoring experiments are conducted as configured in [4, 6]. The maximum length of TRFs is $m = 82$, which is equal to the maximum length of the training sentences. The other configurations are: $K = 300$, $\beta_\lambda = 0.8$, $\beta_\zeta = 0.6$, $t_c = 3000$, $t_0 = 2000$, $t_{max} = 20,000$. L2 regularization with constant 4×10^{-5} is used to avoid over-fitting. 6 CPU cores are used to parallelize the algorithm. The word error rates (WERs) and perplexities (PPLs) on WSJ'92 test set are shown in Tab.2.

Combining TRF and KN5 provides no WER reduction. Different schemes give close WERs for combining TRF and RNN. The log-linear interpolation performs more stable when considering both English and Chinese experiments (as shown later). For English, the obtained WER 7.57% indicates 12.1% and 3.6% relative reductions, when compared to the result of using KN6 (8.61%) and the best result of combining RNN and KN5 (7.85%) respectively.

3.2. Chinese speech recognition on Toshiba dataset

In this section we report the results from using TRF LMs in a large vocabulary Mandarin speech recognition experiment. Different LMs are evaluated by rescoring 30000-best list from a Toshiba's internal test set (2975 utterances). The oracle character error rate (CER) of the 30000-best lists is 1.61%, which are generated with a DNN-based acoustic model. The LM corpus is from Toshiba, which contains about 20M words. We randomly select 1% from the corpus as the development set and others as the training set. The vocabulary contains 82K words, with one special token <UNK>. The NN LM used here is the feedforward neural network (FNN) LM [2, 7] trained by CSLM toolkit[1]. The number of hidden units is 512 and the projection layer units is 3×128. The TRF models are trained using the feature set "w+c+ws+cs+cpw" with different numbers of classes (200,400,600). The configurations are: $m = 100$, $\beta_\lambda = 0.8$, $\beta_\zeta = 0.6$, $t_c = 1000$, $t_0 = t_{max} = 20000$. The sample number K is increased from 300 until no improvements on the development set are observed, and finally set to be 8000. 20 CPU cores are used to parallelize the algorithm. The CERs and PPLs on the test set are shown in Tab.3.

Model	CER	PPL	#feat (M)
KN5 (cutoff=00000)	4.88	227.92	49
KN6 (cutoff=000000)	4.87	227.76	64
FNN [7]	4.30	206.87	53
TRFs ("w+c+ws+cs+cpw")			
200 classes	4.52	∼207	47
400 classes*	4.32	∼191	58
600 classes	4.35	∼186	61
Model Combination			
FNN + KN6 (W)	4.06		
FNN + KN6 (S)	4.16		
FNN + KN6 (Log)	4.11		
TRF* + FNN (S)	4.09		
TRF* + FNN (Log)	**4.0**		

Table 3: The CERs and PPLs on the test set in Chinese speech recognition. "TRF*" denotes the TRF trained with 400 classes.

The results again demonstrate that the TRF LMs significantly outperform the n-gram LMs and are able to match the NN LM. Log-linear combination of TRF and FNN further reduces the CER. The obtained CER 4.0% indicates 17.9% and 1.5% relative reductions, when compared to the result of using KN6 (4.87%) and the best result of combining FNN and KN6 (4.06%) respectively.

4. Related Work and Discussion

Currently, most attention of LM research is attracted by using neural networks, which has been shown to surpass the classic n-gram LMs. Two basic classes of NN LMs are based on FNN [2] and RNN [3]. Recent extensions involve the use of sum-product networks [8], deep recurrent neural networks [9] and feedforward sequential memory networks [10]; only perplexity results are reported in these studies. Crucially, no matter what form the networks take, various NN LMs follow the conditional approach and thus suffer from the expensive softmax computations due to the requirement of local normalization. Lots of studies aim to alleviate this deficiency. Initial efforts include using hierarchical output layer structure with word clustering [3], converting NNs to n-gram LMs [11]. Recently a number of studies [12, 13, 14, 15] make use of noise contrastive estimation (NCE) [16] to build unnormalized variants of NN LMs through trickily avoiding local normalization in training and heuristically fixing the normalizing term in testing.

[1] `http://www-lium.univ-lemans.fr/cslm/`

In contrast, TRF LMs eliminate local normalization from the root and thus are much more efficient in testing with theoretical guarantee. Empirically it is found in [4] that the average time costs for re-ranking of the 1000-best list for a sentence are 0.16 sec vs. 40 sec, based on TRF and RNN respectively (no GPU used). Equally importantly, evaluations in this paper and also in [4, 6] have shown that TRF LMs are able to perform as good as NN LMs (either RNN or FNN) on a variety of tasks. Encouragingly, the random field approach may open a new door to language modeling in addition to the dominant conditional approach, as once envisioned in [5]. Integrating richer features and introducing hidden variables are worthwhile future works.

5. References

[1] S. F. Chen and J. Goodman, "An empirical study of smoothing techniques for language modeling," *Computer Speech & Language*, vol. 13, pp. 359–394, 1999.

[2] H. Schwenk, "Continuous space language models," *Computer Speech & Language*, vol. 21, pp. 492–518, 2007.

[3] T. Mikolov, S. Kombrink, L. Burget, J. H. Cernocky, and S. Khudanpur, "Extensions of recurrent neural network language model," in *Proc. ICASSP*, 2011.

[4] B. Wang, Z. Ou, and Z. Tan, "Trans-dimensional random fields for language modeling," in *Proceedings of the Association for Computational Linguistics (ACL)*, 2015.

[5] R. Rosenfeld, S. F. Chen, and X. Zhu, "Whole-sentence exponential language models: a vehicle for linguistic-statistical integration," *Computer Speech & Language*, vol. 15, pp. 55–73, 2001.

[6] B. Wang, Z. Ou, and Z. Tan, "Learning trans-dimensional random fields with applications to language modeling," *Draft*, 2016.

[7] H. Schwenk, "Cslm - a modular open-source continuous space language modeling toolkit." in *Proc. INTERSPEECH*, 2013, pp. 1198–1202.

[8] W.-C. Cheng, S. Kok, H. V. Pham, H. L. Chieu, and K. M. A. Chai, "Language modeling with sum-product networks." in *INTERSPEECH*, 2014.

[9] W. Zaremba, I. Sutskever, and O. Vinyals, "Recurrent neural network regularization," *arXiv preprint arXiv:1409.2329*, 2014.

[10] S. Zhang, C. Liu, H. Jiang, S. Wei, L. Dai, and Y. Hu, "Feedforward sequential memory networks: A new structure to learn long-term dependency," *arXiv preprint arXiv:1512.08301*, 2015.

[11] E. Arisoy, S. F. Chen, B. Ramabhadran, and A. Sethy, "Converting neural network language models into back-off language models for efficient decoding in automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 1, pp. 184–192, 2014.

[12] A. Mnih and Y. W. Teh, "A fast and simple algorithm for training neural probabilistic language models," *arXiv preprint arXiv:1206.6426*, 2012.

[13] A. Vaswani, Y. Zhao, V. Fossum, and D. Chiang, "Decoding with large-scale neural language models improves translation." in *EMNLP*, 2013, pp. 1387–1392.

[14] A. Sethy, S. Chen, E. Arisoy, and B. Ramabhadran, "Unnormalized exponential and neural network language models," in *Proc. ICASSP*, 2015.

[15] X. Chen, X. Liu, M. Gales, and P. Woodland, "Recurrent neural network language model training with noise contrastive estimation for speech recognition," in *Proc. ICASSP*, 2015, pp. 5411–5415.

[16] M. U. Gutmann and A. Hyvärinen, "Noise-contrastive estimation of unnormalized statistical models, with application to natural image statistics," *The journal of Machine Learning Research*, vol. 13, no. 1, pp. 307–361, 2012.